

- 1 In the claims:
- 2 1. A method for dynamic retransmission of transactions in a multi-processor
3 computer architecture, comprising:
4 at a source node in the computer architecture, the source node comprising a
5 retransmit buffer, designating a transaction for transmission to a destination node in the
6 computer architecture, the destination node comprising a receive buffer, wherein the
7 transaction is designated for transmission over a first path in a first flow control class;
8 retrieving a transaction identification (TID) for the designated transaction;
9 comparing the retrieved TID for the designated transaction to TIDs in the
10 retransmit buffer, wherein if the comparison does not show a match:
11 attaching the retrieved TID to the designated transaction,
12 placing the designated transaction in the retransmit buffer, and
13 sending the designated transaction to the destination node; and
14 wherein if the comparison shows a match, transmitting the designated transaction over a
15 second path.
- 16 2. The method of claim 1, wherein the designated transaction in the retransmit buffer
17 times out, comprising:
18 retrieving a TID of a most recent transaction sent to the destination node along the
19 first path;
20 sending a probe request to the destination node along the second path, the probe
21 request including a TID of the timed-out transaction and the TID of the most recent
22 transaction;
23 deconfiguring the first path; and
24 updating the TID in the retransmit buffer.
- 25 3. The method of claim 2, wherein the destination node receives the probe request,
26 the method at the destination node, comprising:
27 determining a TID for a most recent transaction entry in the receive buffer for a
28 transaction from the source node along the first path and the first flow control class; and
29 determining:
30 if the TID for the most recent transaction entry in the receive buffer equals
31 the timed-out transaction TID.
- 32 4. The method of claim 3, wherein the TIDs are not equal, further comprising
33 designating an error condition.
- 34 5. The method of claim 3, wherein the TID are equal, further comprising:

1 sending a probe response to the source node along the second path;
2 indicating to the source node if the timed-out transaction is received at the
3 destination node; and
4 sending the source node the most recent TID acknowledged.

5 6. The method of claim 5, wherein the source node receives the probe response,
6 comprising resuming transmission for all transactions in the retransmit buffer for which
7 an acknowledgement has not been received.

8 7. The method of claim 5, further comprising:
9 sending a plunge transaction to the destination node over the first path, the plunge
10 transaction indicating a TID for retransmission of transactions should the first path be
11 reconfigured; and
12 updating the TID.

13 8. The method of claim 7, wherein the destination node receives the plunge
14 transaction, the method at the destination node, comprising:
15 determining if the first path is deconfigured from the source node; and
16 if the first path is deconfigured from the source node, sending a plunge response
17 to the source node over the first path.

18 9. The method of claim 8, wherein the destination node determines that the first path
19 is not deconfigured, wherein an error condition exists, comprising:

20 informing the source node; and
21 deconfiguring the first path from the destination node.

22 10. The method of claim 8, further comprising:
23 receiving the plunge response; and
24 reconfiguring the first path from the source node.

25 11. The method of claim 2, wherein the probe request times out, comprising:
26 determining if a third path is available;
27 if the third path is available:

28 sending a second probe response along the third path,
29 deconfiguring the second path, and
30 updating the TID in the retransmit buffer; and

31 if the third path is not available, designating an error condition.

32 12. The method of claim 1, wherein the destination node receives the designated
33 transaction over the first path, the method at the destination node, comprising:
34 determining if the first path is configured;

1 determining from the receive buffer a TID from a most recent transaction from the
2 source node along the first path in the first flow control class;
3 comparing the TIDs of the designated transaction and the most recent transaction
4 to determine if the most recent transaction is in default; and
5 if the most recent transaction is not in default, adding an entry to the receive_TID
6 table.

7 13. The method of claim 12, wherein the most recent transaction is in default,
8 comprising dropping the transaction.

9 14. The method of claim 13, wherein the first path is configured, comprising:
10 accepting the transaction;
11 entering the TID of the transaction receive_TID table; and
12 sending an acknowledgement to the source node along the first path.

13 15. The method of claim 14, further comprising:
14 waiting for a time period slightly less than N times the time of flight, wherein N
15 equals 3 to 4; and
16 invalidating the entry in the receive_TID table.

17 16. An apparatus for retransmission of transactions in a multi-processor computer
18 architecture, comprising:
19 a source node having a retransmit buffer, wherein the source node stores
20 transactions transmitted from the source node;
21 a send_TID table comprising a transaction identification (TID) for each
22 transaction sent from the source node;
23 a destination node comprising a receive buffer, wherein the destination node
24 stores transactions transmitted from the destination node;
25 a receive_TID table comprising a TID for each transaction sent from the
26 destination node, wherein the source node sends normal transactions to the destination
27 node and the destination node sends acknowledgements for the normal transactions to the
28 source node, wherein if the source node does not receive an acknowledgement within a
29 specified time, a corresponding normal transaction in the retransmit buffer times out; and
30 a probe transaction, whereby the source node queries the destination node for a
31 timed-out transaction.

32 17. The apparatus of claim 16, wherein the probe transaction comprises a TID of the
33 timed-out transaction and a TID of a TID of a last transaction pending in the retransmit
34 buffer, wherein the source node comprises:

- 1 means for deconfiguring the first path; and
2 means for updating the TID in the receive_TID table.
- 3 18. The apparatus of claim 17, wherein the destination node receives the probe
4 transaction, the destination node, further comprising:
5 means for determining a TID for a most recent transaction entry in the receive
6 buffer; and
7 means for determining if the TID for the most recent transaction entry in the
8 receive buffer equals the timed-out transaction TID.
- 9 19. The apparatus of claim 18, wherein if either (a) or (b) are met, the destination
10 node:
11 sends a probe response to the source node along the second path;
12 indicates to the source node if the timed-out transaction is received at the
13 destination node; and
14 sends the source node the most recent TID acknowledged.
- 15 20. The apparatus of claim 19, wherein the source node receives the probe response,
16 wherein the source node comprises means for resuming transmission for all transactions
17 in the retransmit buffer for which an acknowledgement has not been received.
- 18 21. The apparatus of claim 18, wherein the TIDs are not equal, further comprising
19 means for designating an error condition.

20200105225007